Partial Quantifier Elimination

Eugene Goldberg and Panagiotis Manolios
Northeastern University, USA, {eigold,pete}@ccs.neu.edu

arXiv:1407.4835v1 [cs.LO] 17 Jul 2014

***Abstract*—We consider the problem of Partial Quantifier Elimination (PQE). Given formula $\exists X[F(X,Y) \wedge G(X,Y)]$, where F, G are in conjunctive normal form, the PQE problem is to find a formula $F^*(Y)$ such that $F^* \wedge \exists X[G] \equiv \exists X[F \wedge G]$. We solve the PQE problem by generating and adding to F^* clauses over the free variables that make the clauses of F with quantified variables *redundant*. The traditional Quantifier Elimination problem (QE) is a special case of PQE where G is empty so *all* clauses of the input formula with quantified variables need to be made redundant. The importance of PQE is twofold. First, many problems are more naturally formulated in terms of PQE rather than QE. Second, in many cases PQE can be solved more efficiently than QE. We describe a PQE algorithm based on the machinery of dependency sequents and give experimental results showing the promise of PQE.**

I. Introduction

The elimination of existential quantifiers is an important problem arising in many practical applications. We will refer to this problem as the Quantifier Elimination problem, or QE. Given a formula $\exists X[F]$ where F is a propositional formula, the **QE problem** is to find a quantifier free formula G such that $G \equiv \exists X[F]$. In this paper, we assume that all propositional formulas are represented in conjunctive normal form (CNF).

Unfortunately, the efficiency of current QE algorithms still leaves much to be desired. This is one reason that many successful theorem proving methods such as interpolation and IC3 avoid QE and use SAT-based reasoning instead. These methods can be viewed as solving specialized versions of the QE problem that can be solved efficiently. For example, finding an interpolant $I(Y)$ of formula $A(X,Y) \wedge B(Y,Z)$ comes down to solving a special case of QE where $I \equiv \exists X[A]$ needs to hold only in subspaces where $B \equiv 1$. So it is important to perform a systematic study of the QE problem, looking for variants of the problem that can be solved efficiently. Such a study can help us better understand existing algorithms that sidestep the use of QE in favor for more limited, specialized methods. The study may also lead to the discovery of new applications of QE.

In this paper, we consider a variation of the QE problem called Partial QE (PQE). Let $\exists X[F(X,Y) \wedge G(X,Y)]$ be a formula where variables of X are quantified. The **PQE problem** is to find a formula $F^*(Y)$ such that $F^* \wedge \exists X[G] \equiv \exists X[F \wedge G]$. We will say that F^* is obtained by **taking *F* out of the scope of the quantifiers**. Note that if $F^* \rightarrow \exists X[G]$ holds, then $F^* \equiv \exists X[F \wedge G]$. That is, in this case, a solution to the PQE problem is also a solution to the QE problem. We will say that in this case QE reduces to PQE.

Our motivation for solving the PQE problem is twofold. First, in many cases, a verification problem can be formulated as an instance of PQE rather than QE. Besides, even if the original problem is formulated in terms QE it can sometimes be reduced to PQE. Second, in many cases, the PQE problem can be solved much more efficiently than QE. We are especially interested in applying PQE when formula F is much smaller than G.

The relation between efficiency of solving PQE and QE can be better understood in terms of clause redundancy [9]. The PQE problem specified by $\exists X[F \wedge G]$ reduces to finding a set of clauses F^* that makes all X-clauses of F redundant in formula $\exists X[F \wedge G]$. (An ***X*-clause** is a clause that contains a variable from X.) Then every clause of F can be either dropped as redundant or removed from the scope of the quantifiers as it contains only free variables.

One can view the process of building F^* as follows. X-clauses of F are made redundant in $\exists X[F \wedge G]$ by adding to F resolvent clauses derived from $F \wedge G$. Notice that no clause obtained by resolving *only* clauses of G needs to be made redundant. Adding resolvents to F goes on until all X-clauses of the current formula F are redundant. At this point, the X-clauses of F can be dropped and the remaining clauses of F form F^*.

If F is much smaller than G, the process of solving PQE looks like wave propagation where F is the original "perturbation" and G is the "media" where this wave propagates. Such propagation can be efficient even if G is large. By contrast, when solving the QE problem for $\exists X[F \wedge G]$ one needs to make redundant the X-clauses of both F and G and *all* resolvent X-clauses including the ones obtained by resolving only clauses of G.

In this paper, we describe a PQE-algorithm called *DS-PQE* that is based on the machinery of D-Sequents [8], [9]. One needs this machinery for PQE for the same reason as for QE [8]. Every clause of $F^*(Y)$ can be obtained by resolving clauses of $F \wedge G$. However, the number of clauses that are implied by $F \wedge G$ and depend only on Y is, in general, exponential in $|Y|$. So it is crucial to identify the moment when the set of clauses derived so far that depend only on Y is sufficient to make the X-clauses of F redundant in $\exists X[F \wedge G]$. The machinery of D-sequents is used for such identification. Namely, one can stop generating new clauses when a D-sequent stating redundancy of the X-clauses of F is derived. We experimentally compare *DS-PQE* with our QE algorithm from [9] in the context of model checking.

The following exposition is structured as follows. In Sections II and III, we discuss some problems that can benefit from an efficient PQE-algorithm. A run of *DS-PQE* on a simple formula is described in Section IV. Sections V and VI

give basic definitions and recall the notion of D-Sequents. In Section VII, *DS-PQE* is described. We discuss previous work in Section VIII. Experimental results are given in Section IX. Finally, we make conclusions in Section X.

II. Using PQE For Model Checking

In this section and the one that follows we describe some applications where using an efficient PQE solver can be very beneficial. A few more applications of PQE are listed in an extended abstract [10].

A. Computing pre-image in backward model checking

Let $T(S, S')$ be a transition relation where S and S' specify the current and next state variables respectively. We will refer to complete assignments $\boldsymbol{s}$ and $\boldsymbol{s'}$ to variables S and S' as present and next states respectively. Let formula $H(S')$ specify a set of next-states and $G(S)$ specify the pre-image of $H(S')$. That is, a present state $\boldsymbol{s}$ satisfies G iff there exists a next state $\boldsymbol{s'}$ such that $H(\boldsymbol{s'}) \wedge T(\boldsymbol{s}, \boldsymbol{s'}) = 1$.

Finding G reduces to QE that is to building a formula logically equivalent to $\exists S'[H \wedge T]$. However, one can construct the pre-image of H by PQE as follows. Let H^* be a formula such that $H^* \wedge \exists S'[T] \equiv \exists S'[H \wedge T]$ *i.e.*, H^* is a solution to the PQE problem. Notice that H^* implies $\exists S'[T]$ because $\exists S'[T] \equiv 1$. Indeed, for every present state $\boldsymbol{s}$ there always exists some next state $\boldsymbol{s'}$ such that $T(\boldsymbol{s}, \boldsymbol{s'}) = 1$. So $H^* \equiv \exists S'[H \wedge T]$ and hence specifies the pre-image of H. In other words, here QE reduces to PQE.

B. State elimination in IC3-like model checkers

In this subsection, we discuss state elimination, a key problem for IC3-like model checkers [2]. Given a transition relation $T(S, S')$, the problem of eliminating a state $\boldsymbol{s}$ is to find a clause C falsified by $\boldsymbol{s}$ and inductive relative to a formula F. The latter means that $F \wedge C(S) \wedge T \rightarrow C(S')$.

The performance of IC3 strongly depends on the efficiency of solving the state elimination problem and the quality of inductive clauses generated to solve it. An IC3-like model checker would benefit from an efficient algorithm finding the pre-image of the state $\boldsymbol{s}$ to be eliminated [13]. Finding the pre-image of $\boldsymbol{s}$ can be useful when no inductive clause C eliminates $\boldsymbol{s}$. In this case, IC3 removes some states that satisfy F and from which a direct transition to $\boldsymbol{s}$ is possible. This is done by adding new clauses to F, which eventually leads to appearance of a clause C that is inductive relative to F and eliminates $\boldsymbol{s}$. Finding the best states to remove is crucial for the performance of IC3. The pre-image of $\boldsymbol{s}$ can be very useful to identify such states.

Finding the pre-image of $\boldsymbol{s}$ is a special case of the problem we discussed in Subsection II-A. Let H be the set of unit clauses specifying state $\boldsymbol{s}$ *i.e.*, $\boldsymbol{s}$ satisfies H. Let $G(S)$ be a formula such that $F \wedge G \wedge \exists S'[T] \equiv F \wedge \exists S'[H(S') \wedge T]$. The complete assignments satisfying G specify the pre-image of $\boldsymbol{s}$ "relative" to F. Any clause C inductive relative to F has to be falsified by assignments satisfying $F \wedge G$. The PQE-algorithm we describe in this paper is not efficient enough to be used in the loop of IC3 right away, but this may change soon.

```
SAT_by_PQE(G){
1  while (true) {
2     C := GenClause(G);
3     R := SolvePQE(∃X[C ∧ G]);
   - - - - - - - - - - - -
4     if (R is derived without using C) {
5        G := AddClause(G, R);
6        if (R ≡ 0) return(UNSAT);
7        continue;}
   - - - - - - - - - - - -
8     if (R ≡ 1) {
9        G := G ∪ {C};
10       continue; }
   - - - - - - - - - - - -
// the only possibility left is R ≡ 0
11    if (G → C) return(UNSAT);
12    else return(SAT);
```

Fig. 1. SAT checking by PQE

III. Using PQE For SAT-solving

In this section, we describe a SAT-algorithm based on PQE. (We will refer to this algorithm as *PQE-SAT*.) We also contrast *PQE-SAT* with a SAT-solver based on Conflict Driven Clause Learning (CDCL).

A. High-level view of the algorithm

The pseudocode of *PQE-SAT* is shown in Figure 1. Let $G(X)$ be a CNF formula to be checked for satisfiability. In the main loop, *PQE-SAT* performs the following actions. First, it generates a clause C that is not trivially subsumed by a clause of G (line 2). Then *PQE-SAT* solves an instance of the PQE problem (line 3). Namely, it calls procedure *SolvePQE* to find formula R such that $R \wedge \exists X[G] \equiv \exists X[C \wedge G]$. Depending on the type of formula R returned by *SolvePQE*, *PQE-SAT* either updates G by adding a clause or makes a final decision on whether G is satisfiable (lines 4-12).

SolvePQE returns three kinds of formula R. The actions *PQE-SAT* take for every kind of formula R are separated by the dotted lines in Figure 1. We will refer to a formula R returned by *SolvePQE* as a formula of the *first kind* if it is obtained by resolving only clauses of G (lines 4-7). In this case, R is just a clause that subsumes C. (In particular, R can be equal to C.) On the one hand, the fact that R is derived without using clause C means that R is implied by G. On the other hand, the fact that R subsumes C suggests that C is also implied by G. Thus C is trivially redundant in $\exists X[C \wedge G]$. *PQE-SAT* adds clause R to G. If clause R is empty, then G is obviously unsatisfiable.

If resolution derivation of the formula R returned by *SolvePQE* involves clause C we will refer to R as a formula of the second or third kind. In this case, R is a constant. That is R either has no clauses (formula of the *second kind*, $R \equiv 1$) or it is an empty clause (formula of the *third find*, $R \equiv 0$). Indeed, just derivation of a clause A subsuming C does not mean that C is redundant in $\exists X[C \wedge G]$. The reason is that A is derived using clause C and so A may not be implied by G. On the other hand, if A is not empty (and hence contains variables of X) , it cannot be taken out of the scope of quantifiers.

Actions of *PQE-SAT* when *SolvePQE* returns a formula R of the second kind are shown in lines 8-10. The fact that $R \equiv 1$ means that C is redundant in $\exists X[C \wedge G]$. That is either C is implied by G or C eliminates some (but not all) assignments satisfying G. In either case, $C \wedge G$ is equisatisfiable to G. For that reason *PQE-SAT* adds C to G.

What *PQE-SAT* does when *SolvePQE* returns a formula R of the third kind is shown in lines 11-12. The fact that $R \equiv 0$ means that either G is unsatisfiable or C is falsified by *every* assignment satisfying G. *PQE-SAT* tells these two cases apart by checking if C is implied by G.

B. Difference between PQE-SAT and a CDCL SAT-solver

The difference between *PQE-SAT* and a CDCL SAT-solver is twofold. First, *PQE-SAT* employs non-resolution derivation of clauses. This derivation occurs, when *SolvePQE* returns a formula R of the second kind (i.e. $R \equiv 1$). In contrast to a formula of the first kind, in this case, *SolvePQE* proves that C is redundant in $\exists X[C \wedge G]$ without generation of a clause subsuming C. A simple example of a clause obtained by non-resolution derivation is a blocked clause [18] (see Section IV). Adding clauses obtained by non-resolution derivation allows one to get proofs that are much shorter than those based on pure resolution. For example, in [17] it was shown that extending resolution with a rule allowing to add blocked clauses makes it exponentially more powerful.

The second difference between *PQE-SAT* and a CDCL SAT-solver is in the way they generate a satisfying assignment. When *SolvePQE* returns an empty clause (a formula of the third kind) it checks if G implies C. A counterexample showing that $G \not\rightarrow C$ is also an assignment satisfying G. Checking if $G \rightarrow C$ holds reduces to testing the satisfiability of G in the subspace where C is falsified.

As far as finding a satisfying assignment is concerned, *PQE-SAT* potentially has three advantages over CDCL-solvers. The first advantage is that *PQE-SAT* can derive clauses that eliminate satisfying assignments of G. This is important because the ability of a CDCL-solver to efficiently find a satisfying assignment hinges on its ability to derive short clauses. For example, if a unit clause $\overline{v}$ is derived by a CDCL-solver, it can immediately set v to 0. However, such a clause cannot be derived if formula G has satisfying assignments with $v = 0$ and $v = 1$. The ability of *PQE-SAT* to add clauses removing satisfying assignments in general leads to enhancing the quality of learned clauses. Suppose, for example, that *PQE-SAT* adds to G a clause C that eliminates all satisfying assignments with $v = 1$ (but preserves at least one satisfying assignment with $v = 0$). Then formula G implies clause $\overline{v}$ and hence the latter can be derived from G by resolution.

The second advantage of *PQE-SAT* is that if clause C is long (i.e. C has many literals), then checking $G \rightarrow C$ can be much simpler than just testing the satisfiability of G. The third advantage of *PQE-SAT* is that in case C is short *PQE-SAT* can exploit the resolution derivation of an empty clause it obtained. Let P denote such a derivation produced by *PQE-SAT*. The fact that G is satisfiable and $C \wedge G$ is not means that every assignment satisfying G falsifies C. This entails that every cut of P must contain either clause C itself or a descendant clause A of C such that $G \not\rightarrow A$. Note that even if C is a short clause, it can have descendants that are very long. So if C is short, one can replace computationally hard check $G \rightarrow C$ with a sequence of checks $G \rightarrow A$ starting with the longest descendant clauses of C.

IV. Example

In this section, we describe a run of a PQE algorithm called *DS-PQE* that is described in Section VII. *DS-PQE* is a modification of the QE algorithm called *DCDS* [9] based on the machinery of Dependency sequents (D-sequents). In this section, we use notions (*e.g.*, that of D-sequents) that will be formally defined in Section VI. Recall that an X-clause is a clause that contains at least one variable from a set X of Boolean variables.

Let $F = C_1 \wedge C_2$ where $C_1 = y \vee x_1$, $C_2 = \overline{y} \vee x_3$ Let $G = C_3 \wedge C_4 \wedge C_5 \wedge C_6$ where $C_3 = \overline{x}_1 \vee x_2$, $C_4 = \overline{x}_1 \vee \overline{x}_2$, $C_5 = \overline{x}_3 \vee x_4$, $C_6 = y \vee \overline{x}_4$. Let $X = \{x_1, x_2, x_3, x_4\}$ be the set of variables quantified in formula $\exists X[F \wedge G]$. So y is the only free variable of $F \wedge G$.

Problem formulation. Suppose one needs to solve the PQE problem of taking F out of the scope of the quantifiers in $\exists X[F \wedge G]$. That is one needs to find $F^*(y)$ such that $F^* \wedge \exists X[G] \equiv \exists X[F \wedge G]$. Below, we describe a run of *DS-PQE* when solving this problem.

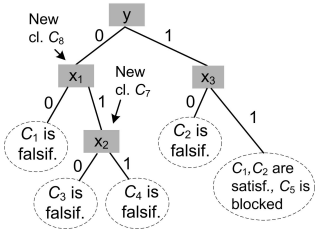

Fig. 2. The search tree built by *DS-PQE*

Search tree. DS-PQE is a branching algorithm. It first proves redundancy of X-clauses of F in subspaces and then merges results of different branches. When *DS-PQE* returns to the root of the search tree, all the X-clauses of F are proved redundant in $\exists X[F \wedge G]$. The search tree built by *DS-PQE* is given in Figure 2. It also shows the nodes where new clauses C_7 and C_8 were derived. *DS-PQE* assigns free variables *before* quantified. For that reason, variable y is assigned first. At every node of the search tree specified by assignment $\boldsymbol{q}$, *DS-PQE* maintains a set of clauses denoted as $PR(\boldsymbol{q})$. Here PR stands for "clauses to Prove Redundant". We will refer to a clause of $PR(\boldsymbol{q})$ as a **PR-clause**. $PR(\boldsymbol{q})$ includes all X-clauses of F plus some X-clauses of G. The latter are proved redundant to make proving redundancy of X-clauses of F easier. Sets $PR(\boldsymbol{q})$ are shown in Figure 4. For every non-leaf node of the search tree two sets of PR-clauses are shown. The set on the left side (respectively right side) of node $\boldsymbol{q}$ gives $PR(\boldsymbol{q})$ when visiting node $\boldsymbol{q}$ for the first time (respectively when backtracking to the right branch of node $\boldsymbol{q}$).

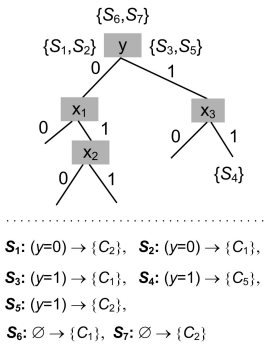

Fig. 3. Derived D-sequents

Using D-sequents. The main concern of *DS-PQE* is to prove redundancy of PR-clauses. Branching is used to reach subspaces where proving redundancy is easy. The redundancy of a PR-clause C is expressed by a Dependency Sequent D-sequent. In short notation, a D-sequent is a record $\boldsymbol{s} \rightarrow \{C\}$ saying that clause C is redundant in formula $\exists X[F \wedge G]$ in any subspace where assignment $\boldsymbol{s}$ is made. We will refer to $\boldsymbol{s}$ as the **conditional part** of the D-sequent. The D-sequents $S_1, \ldots, S_7$ derived by *DS-PQE* are shown in Figure 3. They are numbered in the order they were generated. So-called atomic D-sequents record trivial cases of redundancy. More complex D-sequents are derived by a resolution-like operation called *join*. When *DS-PQE* returns to the root, it derives D-sequents stating the unconditional redundancy of the X-clauses of F.

Merging results of different branches. Let v be the current branching variable and $v = 0$ be the first branch explored by *DS-PQE*. After completing this branch, *DS-PQE* proves redundancy of all clauses that currently have the PR-status. (The only exception is the case when a PR-clause gets falsified in branch $v = 0$. We discuss this exception below.) Then *DS-PQE* explores branch $v = 1$ and derives D-sequents stating redundancy of clauses in this branch. Before backtracking from node v, *DS-PQE* uses operation *join* to produce D-sequents whose conditional part does not depend on v. For example, in branch $y = 0$, D-sequent S_1 equal to $(y = 0) \rightarrow \{C_2\}$ was derived. In branch $y = 1$, D-sequent S_5 equal to $(y = 1) \rightarrow \{C_2\}$ was derived. By joining S_1 and S_5 at variable y, D-sequent S_7 equal to $\emptyset \rightarrow \{C_2\}$ was produced where the conditional part did not depend on y.

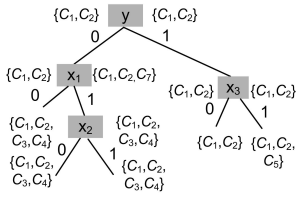

Fig. 4. Dynamics of the $PR(\boldsymbol{q})$ set

Derivation of new clauses. Note that redundancy of the PR-clauses in subspace $y = 1$ was proved without adding any new clauses. On the other hand, proving redundancy of PR-clauses in subspace $y = 0$ required derivation of clauses $C_7 = \overline{x}_1$ and $C_8 = y$. For instance, clause C_7 was generated at node $(y = 0, x_1 = 1)$ by resolving C_3 and C_4. Clause C_7 was *temporarily* added to F to make PR-clauses C_3 and C_4 redundant at the node above. However, C_7 was removed from formula F after derivation of clause C_8 because the former is subsumed by the latter in subspace $y = 0$. This is similar to conflict clause generation in SAT-solvers where the intermediate resolvents are discarded.

Derivation of atomic D-sequents. $S_1, \ldots, S_5$ are the atomic D-sequents derived by *DS-PQE*. They record trivial cases of redundancy. (Due to the simplicity of this example, the conditional part of all atomic D-sequents has only assignment to y *i.e.*, the free variable. In general, however, the conditional part of a D-sequent also contains assignments to quantified variables.) There are three kinds of atomic D-sequents. *D-sequents of the first kind* state redundancy of clauses satisfied in a subspace. For instance, D-sequent S_1 states redundancy of clause C_2 satisfied by assignment $y = 0$. *D-sequents of the second kind* record the fact that a clause is redundant because some other clause is falsified in the current subspace. For instance, D-sequent S_2 states that C_1 is redundant because clause $C_8 = y$ is falsified in subspace $y = 0$. *D-sequents of the third kind* record the fact that a clause is redundant in a subspace because it is blocked at a variable v. That is this clause cannot be resolved on v. For example, D-sequent S_4 states redundancy of C_5 that cannot be resolved on x_4 in subspace $(y = 1, x_3 = 1)$. Clause C_5 is resolvable on x_4 only with C_6 but C_6 is satisfied by assignment $y = 1$.

Computation of the set of PR-clauses. The original set of PR-clauses is equal to the the initial set of X-clauses of F. Denote this set as PR_{init}. In our example, PR_{init} = $\{C_1, C_2\}$. There are two situations where $PR(\boldsymbol{q})$ is extended. The first situation occurs when a parent clause of a new resolvent is in $PR(\boldsymbol{q})$ and this resolvent is an X-clause. Then this resolvent is added to $PR(\boldsymbol{q})$. An example of that is clause $C_7 = \overline{x_1}$ obtained by resolving PR-clauses C_3 and C_4.

The second situation occurs when a PR-clause becomes unit. Suppose a PR-clause C is unit at node $\boldsymbol{q}$, v is the unassigned variable of C and $v \in X$. *DS-PQE* first makes the assignment falsifying C. Suppose that this is assignment $v = 0$. Note that all PR-clauses but C itself are obviously redundant at node $\boldsymbol{q} \cup (v = 0)$. *DS-PQE* backtracks and explores the branch $v = 1$ where clause C is satisfied. At this point *DS-PQE* extends the set $PR(\boldsymbol{q} \cup (v = 1))$ by adding *every* clause of $F \wedge G$ that a) has literal $\overline{v}$; b) is not satisfied; c) is not already in $PR(\boldsymbol{q})$.

The extension of the set of PR-clauses above is done to guarantee that clause C will be proved redundant when backtracking off the node $\boldsymbol{q}$. Let us consider the two possible cases. The first case is that formula $F \wedge G$ is *unsatisfiable* in branch $v = 1$. Then extension of the set of PR-clauses above guarantees that a clause falsified by $\boldsymbol{q} \cup (v = 1)$ will be derived to make the new PR-clauses redundant. Most importantly, this clause will be resolved with C on v to produce a clause rendering C redundant in subspace $\boldsymbol{q}$. The second case is that formula $F \wedge G$ is *satisfiable* in branch $v = 0$. Then the redundancy of the clauses with literal $\overline{v}$ will be proved without derivation of a clause falsified by $\boldsymbol{q} \cup (v = 1)$. When backtracking to node $\boldsymbol{q}$, clause C will be blocked at variable v and hence redundant. Note that extension of the set $PR(\boldsymbol{q})$ is *temporary*. When *DS-PQE* backtracks past node $\boldsymbol{q}$, the clauses that became PR-clauses there lose their PR-status.

Let us get back to our example. The first case above occurs at node $y = 0$ where PR-clause C_1 becomes unit. *DS-PQE* falsifies C_1 in branch $x_1 = 0$, backtracks and explores branch $x_1 = 1$. In this branch, clauses C_3, C_4 of G are made PR-clauses. This branch is unsatisfiable. Making C_3,C_4 PR-

clauses forces *DS-PQE* to derive $C_7 = \overline{x_1}$ that makes C_3, C_4 redundant. But the real goal of obtaining C_7 is to resolve it with C_1 to produce clause $C_8 = y$ that makes C_1 redundant.

The second case above occurs at node $y = 1$ where clause C_2 becomes unit. Clause C_2 gets falsified in branch $x_3 = 0$. Then *DS-PQE* backtracks and explores branch $x_3 = 1$. In this branch, C_5 of G becomes a new PR-clause as containing literal $\overline{x}_3$. This branch is satisfiable and C_5 is proved redundant without adding new clauses. Clause C_2 gets blocked at node $y = 1$ and hence redundant.

Forming a solution to the PQE problem. The D-sequents derived by *DS-PQE* at a node of the search tree are *composable*. This means that the clauses that are redundant individually are also redundant together. For example, on returning to the root node, D-sequents S_6 and S_7 equal to $\emptyset \rightarrow \{C_1\}$ and $\emptyset \rightarrow \{C_2\}$ respectively are derived. The composability of S_6 and S_7 means that D-sequent $\emptyset \rightarrow \{C_1, C_2\}$ holds as well. The only new clause added to F is $C_8 = y$ (clause C_7 was added temporarily). After dropping the X-clauses C_1, C_2 from F as proved redundant one concludes that $y \wedge \exists X[G] \equiv \exists X[F \wedge G]$ and $F^* = y$ is a solution to the PQE problem.

V. Basic Definitions

In this section, we give relevant definitions.

Definition 1: An $\exists$***CNF formula*** is a formula of the form $\exists X[F]$ where F is a Boolean CNF formula, and X is a set of Boolean variables. Let $\boldsymbol{q}$ be an assignment, F be a CNF formula, and C be a clause. $Vars(\boldsymbol{q})$ denotes the variables assigned in $\boldsymbol{q}$; $Vars(F)$ denotes the set of variables of F; $Vars(C)$ denotes the variables of C; and $Vars(\exists X[F]) = Vars(F) \setminus X$.

We consider *true* and *false* as a special kind of clauses.

Definition 2: Let C be a clause, H be a CNF formula, and $\boldsymbol{q}$ be an assignment such that $Vars(\boldsymbol{q}) \subseteq Vars(H)$. Denote by $C_{\boldsymbol{q}}$ the clause equal to *true* if C is satisfied by $\boldsymbol{q}$; otherwise $C_{\boldsymbol{q}}$ is the clause obtained from C by removing all literals falsified by $\boldsymbol{q}$. $H_{\boldsymbol{q}}$ denotes the formula obtained from H by replacing every clause C of H with $C_{\boldsymbol{q}}$. In this paper, we assume that clause $C_{\boldsymbol{q}}$ equal to *true* remains in $H_{\boldsymbol{q}}$. We treat such a clause as *redundant* in $H_{\boldsymbol{q}}$. Let $\exists X[H]$ be an $\exists$CNF and $\boldsymbol{y}$ be an assignment to $Vars(H) \setminus X$. Then $(\exists X[H])_{\boldsymbol{y}} = \exists X[H_{\boldsymbol{y}}]$.

Definition 3: Let S, Q be $\exists$CNF formulas. We say that S, Q are **equivalent**, written $S \equiv Q$, if for all assignments, $\boldsymbol{y}$, such that $Vars(\boldsymbol{y}) \supseteq (Vars(S) \cup Vars(Q))$, we have $S_{\boldsymbol{y}} = Q_{\boldsymbol{y}}$. Notice that $S_{\boldsymbol{y}}$ and $Q_{\boldsymbol{y}}$ have no free variables, so by $S_{\boldsymbol{y}} = Q_{\boldsymbol{y}}$ we mean semantic equivalence.

Definition 4: The **Quantifier Elimination (QE) problem** for $\exists$CNF formula $\exists X[H]$ is to find a CNF formula H^* such that $H^* \equiv \exists X[H]$. The **Partial QE (PQE) problem** for $\exists$CNF formula $\exists X[F \wedge G]$ is to find a CNF formula F^* such that $F^* \wedge \exists X[G] \equiv \exists X[F \wedge G]$.

Definition 5: Let X be a set of Boolean variables, H be a CNF formula and R be a subset of X-clauses of H. The clauses of R are **redundant** in CNF formula H if $H \equiv (H \setminus R)$. The clauses of R are **redundant** in $\exists$CNF formula $\exists X[H]$ if $\exists X[H] \equiv \exists X[H \setminus R]$. Note that $H \equiv (H \setminus R)$ implies $\exists X[H] \equiv \exists X[H \setminus R]$ but the opposite is not true.

VI. Dependency Sequents

In this section, we recall clause Dependency sequents (D-sequents) introduced in [9], operation *join* and the notion of composability. In this paper, we will refer to clause D-sequents as just D-sequents.

Definition 6: Let $\exists X[H]$ be an $\exists$CNF formula. Let $\boldsymbol{s}$ be an assignment to $Vars(H)$ and R be a subset of X-clauses of H. A dependency sequent (**D-sequent**) has the form $(\exists X[H], \boldsymbol{s}) \rightarrow R$. It states that the clauses of $R_{\boldsymbol{s}}$ are redundant in $\exists X[H_{\boldsymbol{s}}]$. Alternatively, we will say that the clauses of R are redundant in $\exists X[H]$ in subspace $\boldsymbol{s}$ (and in any other subspace $\boldsymbol{q}$ such that $\boldsymbol{s} \subseteq \boldsymbol{q}$).

Definition 7: Let $\boldsymbol{s'}$ and $\boldsymbol{s''}$ be assignments in which exactly one variable $v \in Vars(\boldsymbol{s'}) \cap Vars(\boldsymbol{s''})$ is assigned different values. The assignment $\boldsymbol{s}$ consisting of all the assignments of $\boldsymbol{s'}$ and $\boldsymbol{s''}$ but those to v is called the *resolvent* of $\boldsymbol{s'}$,$\boldsymbol{s''}$ on v. Assignments $\boldsymbol{s'}$,$\boldsymbol{s''}$ are called *resolvable* on v.

Definition 8: Let $\exists X[H]$ be an $\exists$CNF formula. Let D-sequents $(\exists X[H], \boldsymbol{s'}) \rightarrow R$ and $(\exists X[H], \boldsymbol{s''}) \rightarrow R$ hold. We refer to these D-sequents as parent ones. Let $\boldsymbol{s'}$, $\boldsymbol{s''}$ be resolvable on $v \in Vars(H)$ and $\boldsymbol{s}$ be the resolvent of $\boldsymbol{s'}$ and $\boldsymbol{s''}$. We will say that D-sequent $(\exists X[H], \boldsymbol{s}) \rightarrow R$ is obtained by **joining** the parents at v. The validity of this D-sequent is implied by that of its parents [9].

Definition 9: Let $\boldsymbol{s'}$ and $\boldsymbol{s''}$ be assignments to a set of variables Z. We will say that $\boldsymbol{s'}$ and $\boldsymbol{s''}$ are *compatible* if every variable of $Vars(\boldsymbol{s'}) \cap Vars(\boldsymbol{s''})$ is assigned the same value in $\boldsymbol{s'}$ and $\boldsymbol{s''}$.

Definition 10: Let $(\exists X[H], \boldsymbol{s'}) \rightarrow R'$ and $(\exists X[H], \boldsymbol{s''}) \rightarrow R''$ be two D-sequents where $\boldsymbol{s'}$ and $\boldsymbol{s''}$ are compatible assignments to $Vars(H)$. We will call these D-sequents **composable** if the D-sequent $(\exists X[H], \boldsymbol{s'} \cup \boldsymbol{s''}) \rightarrow R' \cup R''$ holds.

VII. Algorithm

In this section, we describe a PQE algorithm called ***DS-PQE*** where DS stands for Dependency Sequents. *DS-PQE* is based on our QE algorithm *DCDS* described in [9]. In this section, we will mostly focus on the features of *DS-PQE* that differentiate it from *DCDS*.

DS-PQE derives D-sequents $(\exists X[F \wedge G], \boldsymbol{s}) \rightarrow \{C\}$ stating the redundancy of X-clause C in any subspace $\boldsymbol{q}$ such that $\boldsymbol{s} \subseteq \boldsymbol{q}$. From now on, we will use a short notation of D-sequents writing $\boldsymbol{s} \rightarrow \{C\}$ instead of $(\exists X[F \wedge G], \boldsymbol{s}) \rightarrow \{C\}$. We will assume that the parameter $\exists X[F \wedge G]$ missing in $\boldsymbol{s} \rightarrow \{C\}$ is the *current* $\exists$CNF formula (with all resolvents added to F). One can omit $\exists X[F \wedge G]$ from D-sequents because $(\exists X[F \wedge G], \boldsymbol{s}) \rightarrow \{C\}$ holds no matter how many resolvent clauses are added to F [9]. We will call D-sequent $\boldsymbol{s} \rightarrow \{C\}$ **active** in subspace $\boldsymbol{q}$ if $\boldsymbol{s} \subseteq \boldsymbol{q}$. The fact that $\boldsymbol{s} \rightarrow \{C\}$ is active in subspace $\boldsymbol{q}$ means that C is redundant in $\exists X[F \wedge G]$ in subspace $\boldsymbol{q}$.

// $\boldsymbol{q}$ is an assignment to $Vars(F \wedge G)$
// Ω denotes a set of active D-sequents
// Φ denotes $\exists X[F \wedge G]$
// W denotes $PR(\boldsymbol{q})$
// If *DS_PQE* returns clause *nil* (respectively a non-*nil* clause),
// $(F \wedge G)_{\boldsymbol{q}}$ is satisfiable (respectively unsatisfiable)

DS_PQE (Φ,W,$\boldsymbol{q}$,Ω){
1 if ($\exists$ clause $C \in F \cup G$ falsif. by $\boldsymbol{q}$) {
2 $\Omega := atomic_Dseqs1(\Omega, \boldsymbol{q}, C)$;
3 return(Φ, Ω, C);}
4 $\Omega := atomic_Dseqs2(\Phi, \boldsymbol{q}, \Omega)$;
5* if ($every_PR_clause_redund(W, \Omega)$) return($\Phi, \Omega, nil$);
- - - - - - - - - - - -
6 $v := pick_variable(F \wedge G, \boldsymbol{q}, \Omega)$;
7* (Φ, Ω, C_b) :=*DS_PQE*(Φ,W, $\boldsymbol{q} \cup (v = b)$,$\Omega$);
8 $\Omega_{asym} := Dseqs_to_be_inactive(F, \Omega, v)$;
9 if ($\Omega_{asym} = \emptyset$) return($\Phi, \Omega, C_b$);
10 $\Omega := \Omega \setminus \Omega_{asym}$;
11* if ($impl_assgn(v, \bar{b})$) $W' := new_PR_clauses(W, F \wedge G, \bar{b})$;
12* else $W' := \emptyset$;
13* ($\Phi, \Omega, C_{\bar{b}}$) := *DS_PQE* ($\Phi$,$W \cup W'$,$\boldsymbol{q} \cup (v = \bar{b})$,$\Omega$);
- - - - - - - - - - - - - -
14 if (($C_b \neq nil$) and ($C_{\bar{b}} \neq nil$)){
15 $C := resolve_clauses(C_b, C_{\bar{b}}, v)$;
16 $F := F \wedge C$;
17 $\Omega := atomic_Dseqs1(\Omega, \boldsymbol{q}, C)$;
18* if (($C_b \in W$) or ($C_{\bar{b}} \in W$))
19* $W := W \cup \{C\}$;
20 return(Φ, Ω, C);}
21 $\Omega := merge(\Phi, \boldsymbol{q}, v, \Omega_{asym}, \Omega, C_b, C_{\bar{b}})$;
22 return(Φ, Ω, nil);}

Fig. 5. *DS-PQE* procedure

A. Input and output of DS-PQE

Recall that a PR-clause is an X-clause of $F \wedge G$ whose redundancy needs to be proved in subspace $\boldsymbol{q}$ (see Section IV). A description of *DS-PQE* is given in Figure 5. *DS-PQE* accepts an $\exists$CNF formula $\exists X[F \wedge G]$ (denoted as Φ), an assignment $\boldsymbol{q}$ to $Vars(F)$, the set of PR-clauses (denoted as W) and a set Ω of D-sequents active in subspace $\boldsymbol{q}$ stating redundancy of *some* PR-clauses in $\exists X[F \wedge G]$ in subspace $\boldsymbol{q}$.

Similarly to Section IV, we will assume that the resolvent clauses are added to formula F while formula G remains unchanged. *DS-PQE* returns a formula $\exists X[F \wedge G]$ modified by resolvent clauses added to F (if any), a set Ω of D-sequents active in subspace $\boldsymbol{q}$ that state redundancy of *all* PR-clauses in $\exists X[F \wedge G]$ in subspace $\boldsymbol{q}$ and a clause C. If $(F \wedge G)_{\boldsymbol{q}}$ is unsatisfiable then C is a clause of $F \wedge G$ falsified by $\boldsymbol{q}$. Otherwise, C is equal to *nil* meaning that no clause implied by $F \wedge G$ is falsified by $\boldsymbol{q}$.

The active D-sequents derived by *DS-PQE* are composable. That is if $\boldsymbol{s_1} \rightarrow \{C_1\}, \ldots, \boldsymbol{s_k} \rightarrow \{C_k\}$ are the active D-sequents of subspace $\boldsymbol{q}$, then the D-sequent $\boldsymbol{s^*} \rightarrow \{C_1, \ldots, C_k\}$ holds where $\boldsymbol{s^*} = \boldsymbol{s_1} \cup \ldots \cup \boldsymbol{s_k}$ and $\boldsymbol{s^*} \subseteq \boldsymbol{q}$. Like *DCDS* , *DS-PQE* achieves composability of D-sequents by proving redundancy of PR-clauses in a particular order (that can be different for different paths). This guarantees that no circular reasoning is possible and hence the D-sequents derived at a node of the search tree are composable.

A solution to the PQE problem in subspace $\boldsymbol{q}$ is obtained by discarding the PR-clauses of subspace $\boldsymbol{q}$ (specified by W) from the CNF formula F returned by *DS-PQE*. To solve the original problem of taking F out of the scope of the quantifiers in $\exists X[F \wedge G]$, one needs to call *DS-PQE* with $\boldsymbol{q} = \emptyset$, $\Omega = \emptyset, W = PR_{init}$. Recall that PR_{init} is the set of X-clauses of the original formula F.

B. The big picture

DS-PQE consists of three parts separated in Figure 5 by the dotted lines. In the first part (lines 1-5), *DS-PQE* builds atomic D-sequents recording trivial cases of redundancy of X-clauses. If all the PR-clauses are proved redundant in $\exists X[F \wedge G]$ in subspace $\boldsymbol{q}$, *DS-PQE* terminates at node $\boldsymbol{q}$.

If some PR-clauses are not proved redundant yet, *DS-PQE* enters the second part of the code (lines 6-13). First, *DS-PQE* picks a branching variable v (line 6). Then it recursively calls itself (line 7) starting the left branch of v by adding to $\boldsymbol{q}$ assignment $v = b$, $b \in \{0, 1\}$. Once the left branch is finished, *DS-PQE* explores the right branch $v = \bar{b}$ (line 13).

In the third part, *DS-PQE* merges the left and right branches (lines 14-22). This merging results in proving all PR-clauses redundant in $\exists X[F \wedge G]$ in subspace $\boldsymbol{q}$. For every PR-clause C proved redundant in subspace $\boldsymbol{q}$, the set Ω contains precisely one active D-sequent $\boldsymbol{s} \rightarrow \{C\}$ where $\boldsymbol{s} \subseteq \boldsymbol{q}$. As soon as C is proved redundant, it is marked and ignored until *DS-PQE* enters a subspace $\boldsymbol{q'}$ where $\boldsymbol{s} \not\subseteq \boldsymbol{q'}$ *i.e.*, a subspace where D-sequent $\boldsymbol{s} \rightarrow \{C\}$ becomes inactive. Then clause C gets unmarked signaling that *DS-PQE* does not have a proof of redundancy of C in subspace $\boldsymbol{q'}$ yet.

C. New features of DS-PQE with respect to DCDS

In this paper, we omit the description of functions of Figure 5 that operate identically to those of *DCDS*. What these functions do can be understood from the example of Section IV. If this is not enough, the detailed description of these functions can be found in [9]. In this subsection, we focus on the part of *DS-PQE* that is different from *DCDS*. The code of this part is marked with an asterisk.

The main difference between *DS-PQE* and *DCDS* is that at every node $\boldsymbol{q}$ of the search tree, *DS-PQE* maintains a set $PR(\boldsymbol{q})$ of PR-clauses. $PR(\boldsymbol{q})$ contains all the X-clauses of F and some X-clauses of G (if any). *DS-PQE* terminates its work at node $\boldsymbol{q}$ when all the current PR-clauses are proved redundant (line 5). In contrast to *DS-PQE*, *DCDS* terminates at node $\boldsymbol{q}$, when *all* X-clauses are proved redundant. Line 7 is marked because *DS-PQE* uses an additional parameter W when recursively calling itself to start the left branch of node $\boldsymbol{q}$. Here W specifies the set of PR-clauses to prove redundant in the left branch.

Lines 11-12 show how $PR(\boldsymbol{q})$ is extended. As we discussed in Section IV, this extension takes place when assignment $v = \bar{b}$ satisfies a unit PR-clause C. In this case, the set W' of new PR-clauses is computed. It consists of all the X-clauses that a) contain the literal of v falsified by assignment $v = \bar{b}$; b) are not PR-clauses and c) are not satisfied. As we explained

in Section IV, this is done to facilitate proving redundancy of clause C at node $\boldsymbol{q}$. The set W' is added to W before the right branch is explored (line 13). Notice that the clauses of W' have PR-status only in the subtree rooted at node $\boldsymbol{q}$. Upon return to node $\boldsymbol{q}$ from the right branch, the clauses of W' lose their PR-status.

As we mentioned in Section IV, one more source of new PR-clauses are resolvents (lines 18-19). Let $v = b$ and $v = \overline{b}$ be unsatisfiable branches and C_b and $C_{\overline{b}}$ be the clauses returned by *DS-PQE* . If C_b or $C_{\overline{b}}$ is currently a PR-clause, the resolvent C becomes a new PR-clause. One can think of a PR-clause as supplied with a tag indicating the level up to which this clause preserves its PR-status. If only one of the clauses C_b and $C_{\overline{b}}$ is a PR-clause, then C inherits the tag of this clause. If both parents have the PR-status, the resolvent inherits the tag of the parent clause that preserves its PR-status longer.

D. Correctness of DS-PQE

The correctness of *DS-PQE* is proved similarly to that of *DCDS* [9]. *DS-PQE* is complete because it examines a finite search tree. Here is an informal explanation of why *DS-PQE* is sound. First, the clauses added to F are produced by resolution and so are correct in the sense they are implied by $F \wedge G$. Second, the atomic D-sequents built by *DS-PQE* are correct. Third, new D-sequents produced by operation *join* are correct. Fourth, the D-sequents of individual clauses are composable.

So when *DS-PQE* returns to the root node of the search tree, it derives the correct D-sequent $(\exists X[F \wedge G], \emptyset) \rightarrow F^X$. Here F^X denotes the set of all X-clauses of F. Thus, by removing the X-clauses from F one obtains formula F^* such that $\exists X[F^* \wedge G] \equiv \exists X[F \wedge G]$. Since F^* does not depend on variables of X it can be taken out of the scope of quantifiers.

VIII. Background

QE has been studied by many researchers, due to its important role in verification *e.g.*, in model checking. QE methods are typically based on BDDs [3], [4] or SAT [19], [12], [21], [15], [7], [14], [16]. At the same time, we do not know of research where the PQE problem was solved or even formulated. Of course, identification and removal of redundant clauses is often used in preprocessing procedures of QBF-algorithms and SAT-solvers [6], [1]. However, these procedures typically exploit only situations where clause redundancies are obvious.

PQE is different from QE in at least two aspects. First, a PQE-algorithm has to have a significant degree of "structure-awareness", since PQE is essentially based on the notion of redundancy. So it is not clear, for example, if a BDD-based algorithm would benefit from replacing QE with PQE. This also applies to many SAT-based algorithms of QE. For instance, in [8] we presented a QE algorithm called DDS that was arguably more structure aware than its SAT-based predecessors. DDS is based on the notion of D-sequents defined in terms of variable redundancy. DDS makes quantified variables redundant in subspaces and merges the results of different branches. Despite its structure-awareness, it is hard to adjust DDS to solving PQE: in PQE, one, in general, does not eliminate quantified variables (only some clauses with quantified variables are eliminated).

The second interesting aspect of PQE is as follows. QE can be solved by a trivial albeit inefficient algorithm. Namely, to find a quantifier-free formula equivalent to $\exists X[H]$ one can just resolve out all variables of X as it is done in the DP procedure [5]. However, the PQE problem does not have a counterpart of this algorithm *i.e.*, PQE does no have a "trivial" PQE-solver. Let C be a clause of H and v be a variable of C. One can always make C redundant by adding to H all resolvents of C with clauses of H on v [11], [20]. So one can always "resolve out" any clause of a CNF formula. It seems that one can take formula F out of the scope of the quantifiers in $\exists X[F \wedge G]$ using the following procedure. Keep resolving out clauses of F and their resolvents with G until all non-redundant resolvents depend only on free variables. Unfortunately, this procedure may loop *i.e.*, a previously seen set of clauses $F \wedge G$ may be reproduced later. *DS-PQE* does not have this problem due to branching.

IX. Experimental Results

Since we are not aware of another tool performing PQE, in the experiments we focused on contrasting PQE and QE. Namely, we compared *DS-PQE* with our QE algorithm called *DCDS* [9]. The fact that *DS-PQE* and *DCDS* are close in terms of implementation techniques is beneficial: any difference in performance should be attributed to difference in algorithms rather than implementations.

In the experiments, we used *DS-PQE* and *DCDS* for backward model checking. We will refer to model checkers based on *DS-PQE* and *DCDS* as *MC-PQE* and *MC-QE* respectively. The difference between *MC-PQE* and *MC-QE* is as follows. Let $F(S')$ and $T(S, S')$ specify a set of next-states and transition relation respectively. The basic operation here is to find the pre-image $H(S)$ of F where $H \equiv \exists S'[F \wedge T]$. So H is a solution to the QE problem. As we showed in Subsection II-A, one can also find H just by taking F out of the scope of the quantifiers in formula $\exists S'[F \wedge T]$. *MC-QE* computes H by making redundant *all* S'-clauses of $F \wedge T$ while *MC-PQE* finds H by making redundant *only* the S'-clauses of F.

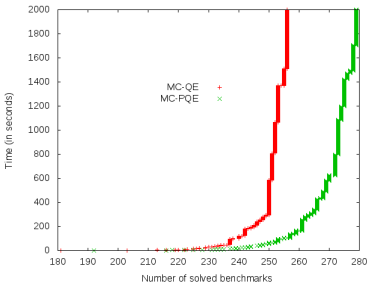

Fig. 6. Performance of model checkers on 282 examples solved by *MC-QE* or *MC-PQE*

The current implementations of *DCDS* and *DS-PQE* lack D-sequent re-using: the parent D-sequents are discarded after a join operation. We believe that re-using D-sequents should boost performance like clause recording in SAT-solving. However, when working on a

new version of *DCDS* we found out that re-using D-sequents indiscriminately may lead to circular reasoning. We have solved this problem theoretically and resumed our work on the new version of *DCDS* . However, here we report the results of implementations that do not re-use D-sequents.

TABLE I
Model checking results on some concrete examples

benchmark	#latches	#gates	#iterations	bug	MC-QE (s.)	MC-PQE (s.)
bj08amba3g62	32	9,825	4	no	241	**38**
kenflashp03	51	3,738	2	no	**33**	104
pdtvishuffman2	55	831	6	yes	>2,000	**296**
pdtvisvsar05	82	2,097	4	no	1,368	**7.7**
pdtvisvsa16a01	188	6,162	2	no	>2,000	**17**
texaspimainp12	239	7,987	4	no	807	**580**
texasparsesysp1	312	11,860	10	yes	39	**25**
pj2002	1,175	15,384	3	no	254	**47**
mentorbm1and	4,344	31,684	2	no	**1.4**	1.7

We compared *MC-PQE* and *MC-QE* on the 758 benchmarks of HWMCC-10 competition [22]. With the time limit of 2,000s, *MC-QE* and *MC-PQE* solved 258 and 279 benchmarks respectively. On the set of 253 benchmarks solved by both model checkers, *MC-PQE* was about 2 times faster (the total time is 4,652s versus 8,528s). However, on the set of 282 benchmarks solved by at least one model checker *MC-PQE* was about 6 times faster (10,652s versus 60,528s). Here we charged 2,000s, *i.e.*, the time limit, for every unsolved benchmark.

Figure 6 gives the performance of *MC-QE* and *MC-PQE* on the 282 benchmarks solved by at least one model checker in terms of the number of problems finished in a given amount of time. Figure 6 shows that *MC-PQE* consistently outperformed *MC-QE* . Model checking results on some concrete benchmarks are given in Table I. The column *iterations* show the number of backward images computed by the algorithms before finding a bug or reaching a fixed point.

ACKNOWLEDGMENT

This research was supported in part by DARPA under AFRL Cooperative Agreement No. FA8750-10-2-0233 and by NSF grants CCF-1117184 and CCF-1319580.

X. CONCLUSION

We introduced the Partial Quantifier Elimination problem (PQE), a generalization of the Quantifier Elimination problem (QE). We presented a PQE-algorithm based on the machinery of D-sequents and gave experimental results showing that PQE can be much more efficient than QE. Efficient PQE-solver may lead to new methods of solving old problems like SAT-solving. In addition, many verification problems can be formulated and solved in terms of PQE rather than QE, a topic ripe for further exploration.

REFERENCES

[1] A.Biere, F.Lonsing, and M.Seidl. Blocked clause elimination for qbf. CADE-11, pages 101–115, 2011.

[2] A. R. Bradley. Sat-based model checking without unrolling. In *VMCAI*, pages 70–87, 2011.

[3] R. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, August 1986.

[4] P. Chauhan, E. M. Clarke, S. Jha, J.H. Kukula, H. Veith, and D. Wang. Using combinatorial optimization methods for quantification scheduling. CHARME-01, pages 293–309, 2001.

[5] M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, July 1960.

[6] N. Eén and A. Biere. Effective preprocessing in sat through variable and clause elimination. In *SAT*, pages 61–75, 2005.

[7] E.Goldberg and P. Manolios. Sat-solving based on boundary point elimination. HVC-10, pages 93–111, 2010.

[8] E.Goldberg and P.Manolios. Quantifier elimination by dependency sequents. In *FMCAD-12*, pages 34–44, 2012.

[9] E.Goldberg and P.Manolios. Quantifier elimination via clause redundancy. In *FMCAD-13*, pages 85–92, 2013.

[10] E.Goldberg and P.Manolios. Software for quantifier elimination in propositional logic. ICMS, 2014. To be published. An extended abstract (four pages in the LNCS format).

[11] A. V. Gelder. Propositional search with k-clause introduction can be polynomially simulated by resolution. In *(Electronic) Proc. 5th Int'l Symposium on Artificial Intelligence and Mathematics*, 1998.

[12] H.Jin and F.Somenzi. Prime clauses for fast enumeration of satisfying assignments to boolean circuits. DAC-05, pages 750–753, 2005.

[13] A. Ivrii. Private communication.

[14] J.Brauer, A.King, and J.Kriener. Existential quantification as incremental sat. CAV-11, pages 191–207, 2011.

[15] J.R.Jiang. Quantifier elimination via functional composition. In *Proceedings of the 21st International Conference on Computer Aided Verification*, CAV-09, pages 383–397, 2009.

[16] W. Klieber, M. Janota, J.M.Silva, and E. M. Clarke. Solving qbf with free variables. In *CP*, pages 415–431, 2013.

[17] O. Kullmann. On a generalization of extended resolution. *Discrete Applied Mathematics*, 34:73–95, 1997.

[18] O. Kullmann. New methods for 3-sat decision and worst-case analysis. *Theor. Comput. Sci.*, 223(1-2):1–72, 1999.

[19] K. McMillan. Applying sat methods in unbounded symbolic model checking. In *Proc. of CAV-02*, pages 250–264. Springer-Verlag, 2002.

[20] M.Järvisalo, M.Heule, and A.Biere. Inprocessing rules. IJCAR-12, pages 355–370, 2012.

[21] M.K.Ganai, A.Gupta, and P.Ashar. Efficient sat-based unbounded symbolic model checking using circuit cofactoring. ICCAD-04, pages 510–517, 2004.

[22] HWMCC-2010 benchmarks, http://fmv.jku.at/hwmcc10/benchmarks.html.